

Joint Probability Estimation Using Tensor Decomposition and Dictionaries

Shaan ul Haque
Department of EE
IIT Bombay
Mumbai, India
shaanhaque2016@gmail.com

Ajit Rajwade
Department of CSE
IIT Bombay
Mumbai, India
ajitvr@cse.iitb.ac.in

Karthik S. Gurumoorthy
Amazon
Bengaluru, India
karthik.gurumoorthy@gmail.com

Abstract—In this work, we study non-parametric estimation of joint probabilities of a given set of discrete and continuous random variables from their (empirically estimated) 2D marginals, under the assumption that the distribution could be approximated by a mixture of product densities/mass functions. Estimation of joint probability density function using semi-parametric techniques such as Gaussian Mixture Models (GMMs) is widely studied. However, they yield poor results when the underlying densities are mixtures of various other families such as Laplacian, generalized Gaussian, uniform, etc. Further, GMMs are not the best choice to estimate distributions which are hybrid in nature, i.e., when it contains both discrete and continuous components. We present a novel approach for estimating the distribution using ideas from dictionary representations in signal processing coupled with low rank tensor decomposition. We create a dictionary of various families of distributions by data inspection, and use it to approximate each decomposed factor of the product in the mixture. Our approach can naturally handle hybrid N -dimensional distributions. We test our approach on a variety of synthetic and real datasets to demonstrate its effectiveness in terms of better classification and lower error rates, when compared to state of the art estimators.

Index Terms—Statistical Learning, Tensor Decomposition, Dictionary Representation, Probability Density Estimation

I. INTRODUCTION

Inferring probability density functions (PDFs) from data is a fundamental problem in machine learning, statistics and signal processing [1] [2] with applications in varied fields such as conditional inference, samples generation, image reconstruction and many more. The applicability of widely popular methods like Gaussian mixture models (GMMs) are restricted to the case of smooth multi-modal densities where every mode is well approximated by Gaussians. Similar is the case with non-parametric settings such as kernel density estimation (KDE). Moreover, such techniques exhibit lower and lower convergence rates as the data dimensionality increases. For N -dimensional data, the convergence rate, in terms of the integrated mean square error (IMSE), for the KDE is known to be $\mathcal{O}(N_s^{-\frac{4}{N+4}})$ [3].

Recently, joint probability mass functions (PMFs) of discrete or discretized random variables (RVs) have been represented as tensors – in fact *low rank* tensors, using the fact that the different RVs are neither completely dependent

nor completely independent [4]. There have been significant developments in the estimation of joint PMFs from lower dimensional marginals (eg, 3D marginals) using the Canonical Polyadic Tensor Decomposition (CPD) using these low rank constraints [4]. Along similar lines, the work in [5] uses non-negative matrix factorization (NMF) techniques to estimate the PMF from just pairwise (i.e. 2D) marginals. On the other hand, ideas from tomography were incorporated into this tensor-based framework in [6] for PMF reconstruction from just 1D marginals. Extending these ideas to the continuous domain, one can discretize the continuous RVs, use the aforementioned techniques for estimating cumulative interval measures (CIMs), followed by an appropriate interpolation technique to recover the joint PDF. For instance, [7] uses sinc interpolation assuming that the underlying PDF is band-limited, in keeping with the popular Shannon-Nyquist theorem. Under similar assumptions, there also exists work in the Fourier domain where CIM reconstruction techniques are applied to obtain the ‘characteristic tensor’ and the continuous PDF is then retrieved using the inverse Fourier Transform [8].

In this work, we present a novel approach which combines ideas from the CPD model for tensors and dictionary representations, to reconstruct the joint PDF from just pairwise (2D) marginals, as opposed to 3D marginals. The key idea is to prepare a dictionary of 1D PDFs belonging to various families, with parameters restricted to lie in a carefully chosen range, for each component of the N -dimensional data. Reconstruction of the N -dimensional PDF using such a dictionary helps us to circumvent restrictive assumptions such as PDF smoothness or band-limitedness as in previous methods. Furthermore, the convergence rate for estimation of 2D marginals is superior to that for 3D marginals used in [4].

II. BACKGROUND

A. Canonical Polyadic Decomposition(CPD) of Tensors

Any N -dimensional tensor $\underline{\mathbf{T}} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$ admits a decomposition in the form of the sum of F rank-1 tensors. This is known as the CPD, and is given by:

$$\underline{\mathbf{T}} = \sum_{r=1}^F \lambda[r] \mathbf{A}_1[:, r] \circ \mathbf{A}_2[:, r] \circ \dots \circ \mathbf{A}_N[:, r], \quad (1)$$

AR wishes to thank SERB Matrix Grant MTR/2019/000691.

where F is the smallest number for which such decomposition is possible, where for each $n \in [N]$, where $[N] \triangleq [1, 2, \dots, N]$ the matrix $\mathbf{A}_n \in \mathbb{R}^{I_n \times F}$ is called a mode factor matrix, \circ denotes the outer product of vectors and $\lambda[r]$ denotes the r^{th} mixing weight. For tensors that represent high-dimensional PMFs, the above decomposition is applicable with the following additional constraint: (1) $\forall n \in [N], r \in [F], \|\mathbf{A}_n[:, r]\|_1 = 1$ with non-negative entries in the mode factor matrices, and (2) $\|\boldsymbol{\lambda}\|_1 = 1$ where $\boldsymbol{\lambda}$ is a vector of F non-negative mixing weights. Recovering the PMF is equivalent to estimating these mode factors and $\boldsymbol{\lambda}$ [4].

The above model can be viewed as a naive Bayes model with the latent variable H that takes on F different values such $P(H = r) = \lambda[r]$ for all $r \in [F]$, and the outer product of the mode factor can be viewed as a conditional probability given H [4]. With this, the CPD model for the PDF of RV $\mathbf{X} = (X_1, X_2, \dots, X_N) \in \mathbb{R}^N$ can be formulated as:

$$\underline{\mathbf{T}} = \sum_{r=1}^F P(H = r) \prod_{n=1}^{n=N} P(X_n = i_n | H = r). \quad (2)$$

B. Continuous RVs: PDF estimation

For the case of continuous RVs, consider an N -dimensional RV $\mathbf{X} = \{X_n\}_{n=1}^N$, whose PDF is given by the following mixture of multivariate distributions:

$$f_{\mathbf{X}}(x_1, \dots, x_N) = \sum_{r=1}^F \lambda[r] f_{\mathbf{X}|H}(x_1, \dots, x_N | H = r). \quad (3)$$

If the RVs are independent given H , then each conditional density can be represented by the product of 1D densities and the above equation becomes:

$$f_{\mathbf{X}}(x_1, \dots, x_N) = \sum_{r=1}^F \lambda[r] \prod_{n=1}^N f_{X_n|H}(x_n | H = r), \quad (4)$$

which can be viewed as the continuous analog of the CPD model for PMFs [7]. In this case, the problem of estimating the PDF is equivalent to estimating these ‘continuous’ mode factors and their mixing weights.

C. Joint PDF estimation from 3D marginals

The work in [7] exploits the above CPD of densities which are conditionally independent to estimate the joint PDF from the data. They propose to discretize each component X_n of the N -dimensional variable into I_n intervals $\{\Delta_n^i \triangleq (d_{n-1}^i, d_n^i)\}_{1 \leq i \leq I_n}$ and form the CIM tensor $\underline{\mathbf{Z}}$ given by:

$$\underline{\mathbf{Z}}(i_1, \dots, i_N) = P(X_1 \in \Delta_1^{i_1}, \dots, X_N \in \Delta_N^{i_N}) \quad (5)$$

$$= \sum_{r=1}^F \lambda[r] \prod_{n=1}^N P(X_n \in \Delta_n^i | H = r), \quad (6)$$

where $P(X_n \in \Delta_n^i | H = r)$ are the CIMs of the corresponding 1D components of the RV. If the PDF/CDF of each component in the product is band-limited, the CIMs can be estimated using the PMF reconstruction techniques described in [4]. The reconstruction involves estimating the 3D marginals,

$\underline{\mathbf{Z}}_{i,j,k} = P(X_i, X_j, X_k)$, using standard histogramming and then minimizing the following cost function:

$$\min_{\{\mathbf{A}_n\}_{n=1}^N, \boldsymbol{\lambda}} \sum_i \sum_{j>i} \sum_{k>j} \|\underline{\mathbf{Z}}_{i,j,k} - [\boldsymbol{\lambda}, \mathbf{A}_i, \mathbf{A}_j, \mathbf{A}_k]\|_F^2$$

$$\text{s.t. } \forall n, r \|\mathbf{A}_n[:, r]\|_1 = \|\boldsymbol{\lambda}\|_1 = 1, \mathbf{A}_n \succeq \mathbf{0}, \boldsymbol{\lambda} \succeq \mathbf{0}, \quad (7)$$

where \succeq represents the element-wise inequality and $[\boldsymbol{\lambda}, \mathbf{A}_i, \mathbf{A}_j, \mathbf{A}_k] \triangleq \sum_{r=1}^F \lambda[r] \mathbf{A}_i[:, r] \circ \mathbf{A}_j[:, r] \circ \mathbf{A}_k[:, r]$. After obtaining the discretized samples of the PDFs, the technique in [7] invokes the Shannon-Nyquist sampling theorem to use sinc interpolation to produce the original joint PDF using Eqn. 1. In [8], under similar band-limitedness of the density, a different kind of approach is considered using the characteristic function of the density, $\Phi_{\mathbf{X}}(\boldsymbol{\nu}) = E[e^{j\boldsymbol{\nu}^T \mathbf{X}}]$, where $j \triangleq \sqrt{-1}$. If the true density is given by $f_{\mathbf{X}}(\mathbf{x})$, then it can be well approximated by truncating the Fourier series below:

$$\hat{f}_{\mathbf{X}}(\mathbf{x}) = \sum_{k_1=-K_1}^{k_1=K_1} \dots \sum_{k_N=-K_N}^{k_N=K_N} \Phi_{\mathbf{X}}(\mathbf{k}) e^{-2\pi j \mathbf{k}^T \mathbf{x}}. \quad (8)$$

If we further impose the CPD model on it, the expression for $\Phi_{\mathbf{X}}(\mathbf{k})$ becomes similar to Eq. 2. where F is the rank of the tensor. With enough samples $\{\mathbf{x}_m\}$, the expectation can be reliably estimated using sample mean $\hat{\Phi}_{\mathbf{X}}(\boldsymbol{\nu}) = \frac{1}{M} \sum_{m=1}^M e^{j\boldsymbol{\nu}^T \mathbf{x}_m}$. Minimizing a cost function similar to Eqn.7 for the characteristic tensor $\hat{\Phi}_{\mathbf{X}}(\boldsymbol{\nu})$ followed by inverse Fourier transform, yields $f_{\mathbf{X}}(\mathbf{x})$.

D. Joint PMF estimation from 2D marginals

An interesting approach that employs Non-negative Matrix Factorization (NMF) techniques [9] to estimate the mode factors from 2D marginals was introduced in [5]. They estimate $\underline{\mathbf{Z}}_{j,k}$ via sample histogramming and obtain the mode factors using the relation $\underline{\mathbf{Z}}_{j,k} = \mathbf{A}_j \boldsymbol{\Lambda} \mathbf{A}_k^T$, where $\boldsymbol{\Lambda}$ is a diagonal matrix with diagonal elements obtained from $\boldsymbol{\lambda}$. If the tensor rank $F \gg \min(I_j, I_k)$, then NMF techniques cannot be applied [10]. Therefore, the authors proposed to split the indices of N variables into two sets and construct a matrix $\tilde{\mathbf{Z}}$ by row and column concatenation using the indices in the two sets (see [5, Eqn. 3]). Then $\tilde{\mathbf{Z}}$ is decomposed as $\tilde{\mathbf{Z}} = \mathbf{W} \mathbf{H}^T$ using the successive projection algorithm (SPA) [11]. The mode factors are then extracted using the relations $\mathbf{W} = [\mathbf{A}_{l_1}, \mathbf{A}_{l_2}, \dots, \mathbf{A}_{l_M}]^T$ and $\mathbf{H}^T = \boldsymbol{\Lambda} [\mathbf{A}_{l_{M+1}}, \mathbf{A}_{l_{M+2}}, \dots, \mathbf{A}_{l_N}]$, where $\{l_1, l_2, \dots, l_M\}$ and $\{l_{M+1}, l_{M+2}, \dots, l_N\}$ are the two sets of indices.

III. PROBLEM STATEMENT AND ALGORITHM

Let $\mathbf{X} \triangleq (X_1, X_2, \dots, X_N)$ be an N -d RV, each component of which can be either continuous or discrete. Our aim is to estimate $f_{\mathbf{X}}(\mathbf{x})$ which follows the CPD model with rank F for the continuous case, given sample values of the RV. Let us further assume that each column of these ‘‘continuous mode factors’’ $f_{X_n|H}(x_n | H = r)$ are convex combinations of various densities from a given dictionary. Mathematically, we have $f_{X_n|H}(x_n | H = r) = \mathcal{A}_n[:, r] = \mathcal{D}_n \mathbf{B}_n[:, r]$, $1 \leq r \leq F$, where \mathcal{D}_n is a dictionary of continuous densities

(or discrete PMFs in some cases) and $\mathbf{B}_n[:, r] \in \mathbb{R}_+^{L_n} \cup \{0\}$ is the non-negative weight vector which sums to 1. Here L_n is the number of different densities (number of columns) that are present in the dictionary \mathcal{D}_n . We will later see that keeping the dictionary \mathcal{D}_n separate for each component of the RV gives us a lot of flexibility in dealing with RVs defined on disparate domains. Then, the 2D PDFs $\mathcal{Z}_{j,k}$ can be expressed in the form $\mathcal{Z}_{j,k} = \mathcal{D}_j \mathbf{B}_j \Lambda \mathbf{B}_k^T \mathcal{D}_k^T$ which is straightforward to derive by marginalizing along X_j and X_k , and replacing each density in the product by its dictionary representation. We discretize each component of the RV into I_n intervals $\{\Delta_n^i \triangleq (d_n^{i-1}, d_n^i)\}_{1 \leq i \leq I_n}$ and form the 2D PMF matrix $\tilde{\mathcal{Z}}_{j,k}$ given by:

$$\begin{aligned} \tilde{\mathcal{Z}}_{j,k}(i_j, i_k) &= P(X_j \in \Delta_j^{i_j}, X_k \in \Delta_k^{i_k}) \\ &= \bar{\mathbf{D}}_j[i_j, :] \mathbf{B}_j \Lambda \mathbf{B}_k^T \bar{\mathbf{D}}_k[i_k, :]^T, \end{aligned} \quad (9)$$

where the column-wise discretized form of a ‘‘continuous tensor/matrix’’ \mathcal{Z} is represented as $\tilde{\mathcal{Z}}$. We can obtain an estimate $\hat{\mathcal{Z}}_{j,k}$ of these 2D PMFs by standard histogramming of the samples. It is emphasised here that we do not assume any knowledge of the densities from which the data are generated. For the purpose of estimation, all the dictionaries $\{\mathcal{D}_n\}_{n=1}^N$ are designed by inspecting the samples. We will elaborate on this aspect in more detail in Sec.IV. Thus, our task is to estimate the coefficients $\Lambda, \mathbf{B}_1, \dots, \mathbf{B}_N$. To this end, we minimize the following cost function:

$$J(\{\mathbf{B}_n\}_{n=1}^N, \Lambda) = \sum_{j,k < N} \|\hat{\mathcal{Z}}_{j,k} - \bar{\mathbf{D}}_j \mathbf{B}_j \Lambda \mathbf{B}_k^T \bar{\mathbf{D}}_k^T\|_F^2 \quad (10)$$

$$\text{s.t. } \forall j, r, \|\mathbf{B}_j[:, r]\|_1 = 1, \mathbf{B}_j \succeq \mathbf{0}, \|\text{diag}(\Lambda)\|_1 = 1, \Lambda \succeq \mathbf{0}.$$

However, minimization via (say) a simple gradient descent may not be feasible as the solution will not be identifiable for cases when $F > \min(L_j, L_k)$, as argued in [5]. We propose to do the following: (i) We minimize $\|\hat{\mathcal{Z}}_{j,k} - \bar{\mathbf{D}}_j \mathbf{T}_{j,k} \bar{\mathbf{D}}_k^T\|_F^2$ for each pair of (j, k) via mirror descent to obtain $\mathbf{T}_{j,k} \triangleq \mathbf{B}_j \Lambda \mathbf{B}_k^T$. For mirror descent, closed-form updates can be derived even with the simplex constraint on $\mathbf{T}_{j,k}$ [12] [13]. (ii) Next, we construct the matrix $\tilde{\mathbf{T}}$ from $\mathbf{T}_{j,k}$ using the concatenation approach described in [5, Eqn. 3] and determine the matrices \mathbf{W} and \mathbf{H}^T as outputs from function SPA(.) implemented in [5]. The weight matrices, \mathbf{B}_n and Λ can now be identified as submatrices of \mathbf{W} and \mathbf{H}^T . (iii) We further refine our estimates for these weights by using mirror descent on cost function in Eqn. 10. We name our algorithm **JUPAD: Joint density estimation Using Pairwise marginals And Dictionaries**.

The pseudo-code for the algorithm is presented in Alg.1, where η_T, η_B and η_L are the learning rate hyper-parameters chosen via cross-validation, \otimes represents the Hadamard product of two matrices, and $\text{vec}(\cdot)$ reshapes a matrix into vector form.

Algorithm 1 Joint probability density estimation using pairwise marginals

- 1: **Procedure: JUPAD**
 - 2: Obtain the estimate $\hat{\mathcal{Z}}_{j,k}$ for the 2D marginals $\mathcal{Z}_{j,k}$ via histogramming
 - 3: **for** each pair (j, k) , $j < k$ **do**
 - 4: Randomly initialize $\mathbf{T}_{j,k}$
 - 5: **while** converged==false **do**
 - 6: $\mathbf{T}_{j,k} \leftarrow \mathbf{T}_{j,k} \otimes \exp(-\eta_T \frac{\partial(\|\hat{\mathcal{Z}}_{j,k} - \bar{\mathbf{D}}_j \mathbf{T}_{j,k} \bar{\mathbf{D}}_k^T\|_F^2)}{\partial \mathbf{T}_{j,k}})$
 - 7: $\mathbf{T}_{j,k} \leftarrow \frac{\mathbf{T}_{j,k}}{\|\text{vec}(\mathbf{T}_{j,k})\|_1}$
 - 8: **end while**
 - 9: **end for**
 - 10: Assemble $\tilde{\mathbf{T}}$ from $\{\mathbf{T}_{j,k}\}$ following [7, Eqn. 3].
 - 11: $\{\mathbf{B}_n\}_{n=1}^N, \Lambda \leftarrow \text{SPA}(\tilde{\mathbf{T}})$
 - 12: **while** converged==false **do**
 - 13: **for** i=1 to N **do**
 - 14: **while** converged==false **do**
 - 15: $\mathbf{B}_n \leftarrow \mathbf{B}_n \otimes \exp(-\eta_B \frac{\partial J}{\partial \mathbf{B}_n})$
 - 16: L_1 normalize each column of \mathbf{B}_n
 - 17: **end while**
 - 18: **end for**
 - 19: **while** converged==false **do**
 - 20: $\Lambda \leftarrow \Lambda \otimes \exp(-\eta_L \frac{\partial J}{\partial \Lambda})$
 - 21: L_1 normalize the diagonal of Λ
 - 22: **end while**
 - 23: **end while**
-

IV. NUMERICAL RESULTS

A. Synthetic Data

To test our algorithm, we created density functions as a convex combinations of densities from a chosen dictionary. The aim then was to reconstruct the density function from samples of the underlying random variable. We drew sample data from the synthetic density for various sample sizes (N_s) and tested the accuracy of the algorithms by averaging the absolute of log likelihood ratio between the estimated ($\hat{f}_{\mathbf{X}}(\cdot)$) and the known true PDF ($f_{\mathbf{X}}(\cdot)$). We generated $M \triangleq 1000$ test samples $\{z_k\}_{k=1}^M$ and used the following measure:

$$D(\hat{f}_{\mathbf{X}}, f_{\mathbf{X}}) \approx \frac{1}{M} \sum_{k=1}^M |\log(\hat{f}_{\mathbf{X}}(z_k)/f_{\mathbf{X}}(z_k))|. \quad (11)$$

Dictionary Construction: Notice that the marginalized distribution for the n^{th} component of the RV will be $f_{X_n}(x_n) = \mathcal{D}_n \mathbf{B}_n \lambda$. This implies that by looking at the structure of the 1D empirical histogram (obtained from sample values) for each component, we can guess the families of distributions it might belong to. For example, the histogram of an exponential RV will peak at zero and decrease exponentially, that of a mixture of Laplacians will have distinct peaks at the mean of each component with a heavy tail, etc. For the parameters of the density families, we consider the range of values for samples of that component, say $[a, b]$. Consider that we wish to include Gaussians and Laplacians in our

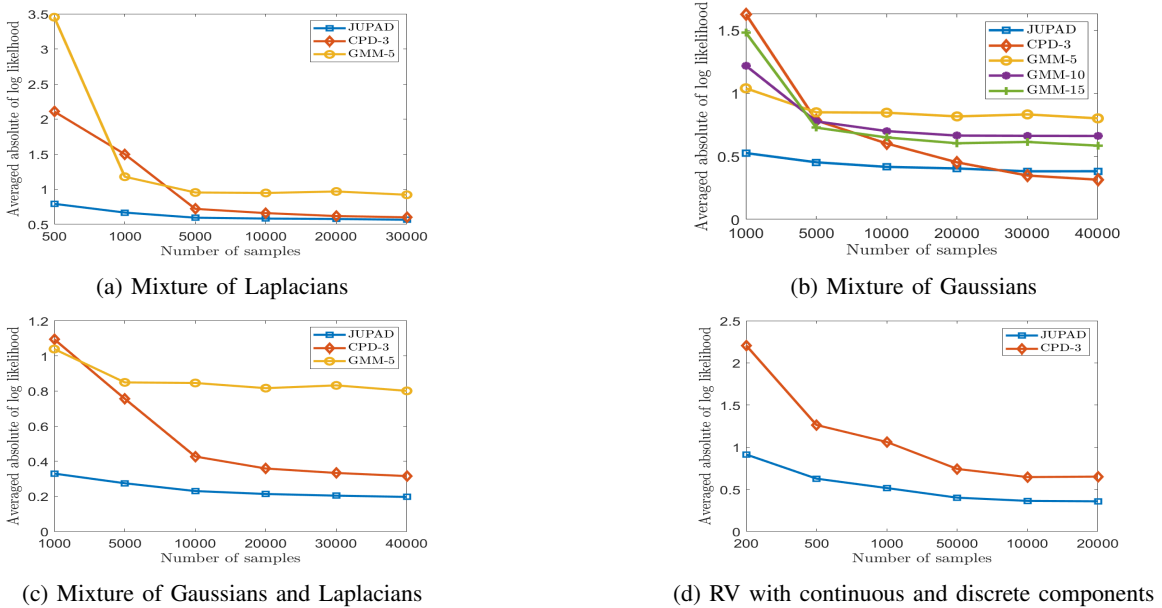


Fig. 1: $D(\hat{f}_{\mathbf{X}}, f_{\mathbf{X}})$ vs Number of samples (N_s) for JUPAD (our approach), CPD-3 [7] and GMM.

dictionaries. Then we can divide $[a, b]$ into regular intervals, and use the interval boundaries as the mean values for those distributions, while the variance/shape factors are chosen such that the densities with the mean values considered as described here, are sufficiently separated. In other words, we want the densities in our dictionary to completely cover the range of the samples for that component. If the n^{th} component of the data is discrete, say label, categorical or integer data with C_n states, we simply assign an identity matrix of size $C_n \times C_n$ to $\bar{\mathbf{D}}_n$. In this case the PMF of the marginal is given by: $f_{X_n}(x_n = i) = \sum_{r=1}^F \lambda[r] \mathbf{B}_n[i, r]$. See appendix of [14] for more details.

For all the experiments we compared our method with the PDF estimation algorithm from [7] and the well-known expectation maximization (EM) algorithm for GMM fitting [15]. For EM-GMM, GMM- n represents the performance plot for a GMM with n clusters. In all experiments, we started with 5 clusters and went up to the point where increasing the number of clusters resulted in ill-conditioned co-variance. We refer to the algorithm in [7] as ‘‘CPD-3’’. The technique described in [8] was not included in the comparison results, as for many toy experiments it produced results similar to that of CPD-3 but was computationally very expensive for larger number of samples. We generated densities belonging to various families, as described below:

Mixture of Laplacians: In the first experiment, we chose the dimension of the RV to be $N = 5$ and $F = 10$. Each column in the mode factor is a mixture of 5 Laplacians. Thus, if $\mathcal{L}(\mu, \alpha)$ denotes a Laplacian density with mean μ and shape factor α , then each $f_{X_n|H}(x_n|H = r) = \sum_{i=1}^5 w_{i,nr} \mathcal{L}(\mu_{i,nr}, \alpha_{i,nr})$ where $\mu_{i,nr} \sim \mathcal{U}(-5, 5)$ and $\alpha_{i,nr} \sim \mathcal{U}(1, 2)$. Here and for all the following cases as well, the mixing weights $w_{i,nr}$ and $\lambda[r] \sim \mathcal{U}(0, 1)$ and then L_1 normalized. Fig. 1a shows

$D(\hat{f}_{\mathbf{X}}, f_{\mathbf{X}})$ vs number of samples (N_s). Clearly, our method works better, be it in low or high sample regime, as Laplacian densities are neither smooth nor band-limited.

Mixture of Gaussians: In the second experiment, we chose $N = 6$ and $F = 8$. Each column in the mode factor was chosen to be a mixture of 5 Gaussian densities. Thus, $f_{X_n|H}(x_n|H = r) = \sum_{i=1}^5 w_{i,nr} \mathcal{N}(\mu_{i,nr}, \sigma_{i,nr}^2)$ where $\mu_{i,nr} \sim \mathcal{U}(-5, 5)$ and $\sigma_{i,nr}^2 \sim \mathcal{U}(1, 2)$. Fig. 1b shows us that in lower sample regime, our algorithm is significantly accurate whereas in high sample regime it is not far from the best method. The reason for this performance anomaly is due to the fact that there is a limit to the accuracy with which the dictionary elements can represent a density due to parameter discretization while the algorithm CPD-3 performs sinc-interpolation which does not have any such restriction.

Mixture of Gaussians and Laplacians: In this experiment, we chose $N = 7$ and $F = 10$. Each column in the first five mode factors is a mixture of 5 Laplacians while mixture of 5 Gaussian densities were used for next 5 mode factors. Again $\mu_{i,nr} \sim \mathcal{U}(-5, 5)$ for both Gaussians and Laplacians, $\alpha_{i,nr}, \sigma_{i,nr}^2 \sim \mathcal{U}(1, 2)$. Just like the first case, our method performs significantly better than other algorithms across all sample regimes, as seen in Fig. 1c.

Mixture of Continuous and Discrete RVs: For the last experiment, we chose $N = 4$ and $F = 8$. This time the last component in the RV was discrete with 10 states. As explained earlier, the dictionary $\bar{\mathbf{D}}_4$ was chosen to be a 10×10 identity matrix. Each column of the last mode factor $\mathbf{B}_4[:, r]$ was generated from $\mathcal{U}(0, 1)$ and then L_1 normalized. Densities of the other mode factors were generated in the same way as for experiment 2. We could not compare our method with the EM-GMM algorithm because of its incapability to incorporate the discrete components of the RV in its formulation. The results




Fig. 2: Classification accuracy on various datasets

for this experiment are presented in Fig. 1d.

B. Real Data

In the real-world datasets where the true underlying PDF is unknown, we tested our algorithm for a classification task on various datasets in UCI repository¹ – ‘Banknote Authentication’ (5D), ‘Wifi Localization’ (8D), ‘Raisins’ (8D) and ‘Seeds’ (8D) datasets, and the KTH TIPS dataset². The classification results for different methods are summarized in Fig. 2. The flexibility of our algorithm to adapt to hybrid distributions and learn the joint probability of both discrete and continuous components, endows us to estimate the joint density of the form $p_{\mathbf{X},Y}(\mathbf{x},y)$, where \mathbf{X} denotes the vector of class feature (continuous) and Y is their label (discrete). As described above, the dictionaries \mathcal{D}_n for the features \mathbf{X} are chosen from the continuous distribution families after examining their empirical marginals. For the label Y , \mathcal{D}_n is set to identity matrix. For detailed information on dictionary construction refer to the appendix of [14]. The value of F was chosen on the basis of accuracy on a validation set which was distinct from the training and test sets. For all the experiments, the classification task was performed by MAP estimation which assigns the label $\hat{y} = \operatorname{argmax}_y p(y|\mathbf{x}) = \operatorname{argmax}_y p(\mathbf{x},y)/p(\mathbf{x})$. We ran the CPD-3 algorithm as implemented by the authors. Treating the label of each class as the latent variable, the classification in CPD-3 is done using the MLE estimate: $\operatorname{argmax}_y p(\mathbf{x}|y)$, assuming the prior $p(y)$ to be uniform, unlike our algorithm where we learn the complete joint density $p(\mathbf{x},y)$. We also compared with GMM having full covariance (‘GMM-Full’) and diagonal covariance matrices (‘GMM-Diag’).

UCI Dataset: Here we would like to bring out the generality of our method to model discontinuities in the PDFs. For eg., if the histogram of some component shows abrupt change at some value, then we can model this discontinuity by keeping few uniform distributions spanning the range of the histogram along with other densities. For most of the data, using a dictionary consisting of Gaussians and Uniform densities yielded satisfactory results (see Fig. 2).

KTH TIPS: This is a texture dataset which contains images of size 200×200 of various textures. We chose three textures

¹<https://archive.ics.uci.edu/ml/datasets.php>

²<https://www.csc.kth.se/cvap/databases/kth-tips/download.html>

– Orange Peel, Bread and Linen for our classification task. We used two training images from each class and divided them into 5×5 patches creating 26D (patch size+label) data. Thereafter, we normalized the pixel values so that all of them lie in the range $[0, 1]$. Here, we used a dictionary consisting of only Gaussians. For testing, we created a collage of 5×5 patches of these textures and classified each patch again by using MAP estimator. Our model performed remarkably well for such a high-dimensional data and outperformed all the other algorithms by a reasonable margin (see Fig. 2).

V. CONCLUSION AND FUTURE WORK

We integrated ideas from low-rank tensors CPD and dictionary representation of signals to present a novel joint density estimation technique. Our method is completely general and can be applied to model mixtures of distributions coming from different families. The numerical results demonstrate the efficacy of our algorithm, especially in the low sample regime where other methods under-perform. Some future work may include a theoretical analysis of the proposed method w.r.t. sample complexity. Our method can also be extended to dictionary learning, where the dictionary elements are themselves learned from the data without any manual inspection.

REFERENCES

- [1] Kevin Patrick Murphy, *Machine Learning: A Probabilistic Perspective*, MIT Press, 2012.
- [2] M. Wainwright, *High-dimensional Statistics: A Non-asymptotic Viewpoint*, vol. 88, 2020.
- [3] Y.-C. Chen, “Lecture 6: Density estimation: Histogram and kernel density estimator,” http://faculty.washington.edu/yenchic/18W_425/Lec6_hist_KDE.pdf.
- [4] N. Kargas, N. D. Sidiropoulos, and X. Fu, “Tensors, learning, and “kolmogorov extension” for finite-alphabet random vectors,” *IEEE Transactions on Signal Processing*, vol. 66, no. 18, 2018.
- [5] S. Ibrahim and X. Fu, “Recovering joint probability of discrete random variables from pairwise marginals,” *IEEE Transactions on Signal Processing*, vol. 69, pp. 4116–4131, 2021.
- [6] J. Vora, K. S. Gurumoorthy, and A. Rajwade, “Recovery of joint probability distribution from one-way marginals: Low rank tensors and random projections,” in *IEEE Statistical Signal Processing*, 2021.
- [7] N. Kargas and N. D. Sidiropoulos, “Learning mixtures of smooth product distributions: Identifiability and algorithm,” in *AISTATS*, 2019.
- [8] M. Amiridi, N. Kargas, and N. D. Sidiropoulos, “Low-rank characteristic tensor density estimation part I: Foundations,” <https://arxiv.org/abs/2008.12315>.
- [9] Daniel Lee and H. Sebastian Seung, “Algorithms for non-negative matrix factorization,” in *Advances in Neural Information Processing Systems*, T. Leen, T. Dietterich, and V. Tresp, Eds. 2000, vol. 13, MIT Press.
- [10] Xiao Fu, Kejun Huang, and Nicholas D. Sidiropoulos, “On identifiability of nonnegative matrix factorization,” *IEEE Signal Processing Letters*, vol. 25, no. 3, pp. 328–332, 2018.
- [11] N. Gillis and S. A. Vavasis, “Fast and robust recursive algorithms for separable nonnegative matrix factorization,” *IEEE TPAMI*, vol. 36, no. 4, pp. 698–714, 2014.
- [12] Jyrki Kivinen and Manfred K. Warmuth, “Exponentiated gradient versus gradient descent for linear predictors,” *Information and Computation*, vol. 132, no. 1, pp. 1–63, 1997.
- [13] Amir Beck and Marc Teboulle, “Mirror descent and nonlinear projected subgradient methods for convex optimization,” *Operations Research Letters*, vol. 31, no. 3, pp. 167–175, 2003.
- [14] S. ul Haque, A. Rajwade, and K. S. Gurumoorthy, “Joint probability estimation using tensor decomposition and dictionaries,” <https://arxiv.org/abs/2203.01667>.
- [15] A. P. Dempster, N. M. Laird, and D. B. Rubin, “Maximum likelihood from incomplete data via the em algorithm,” *JOURNAL OF THE ROYAL STATISTICAL SOCIETY, SERIES B*, vol. 39, no. 1, pp. 1–38, 1977.